Shibboleth Identity Provider Productionalization

SWITCH

Serving Swiss Universities

Chad La Joie
chad.lajoie@switch.ch

What is the Goal?

- An infrastructure in which the mean time between service interruption is very low, response time is acceptable, and staff sanity is maintained.

- What is “very low” and “acceptable”?

Staff

- Train staff in all the software that is used (OS, Java, servlet container, IdP)
- Ensure you have enough people that if you lose 2-3 you don’t lose institutional knowledge.
- Support your staff if they say something is a bad idea.

Hardware

- Multiple servers
  - Enough to handle peak load + 20%
- **Maybe** redundant components within a single piece of hardware
- Beware of blade servers, many don’t have redundant backplanes.

Network

- Hardware/Appliance Load Balancers in either active/standby or active/active configurations
  - Do not use DNS, or other supposed substitutions
- Track your network connections. Ensure they don’t ultimately connect to a single switch or router.

Other Items

- Ensure every authentication system and data source the IdP depends on meets all the previous requirements as well.
- Create a pre-production environment that mirrors the production one (with perhaps a few less servers).
- Monitor logs for WARN and ERROR messages and email/page people.

Good System Architecture

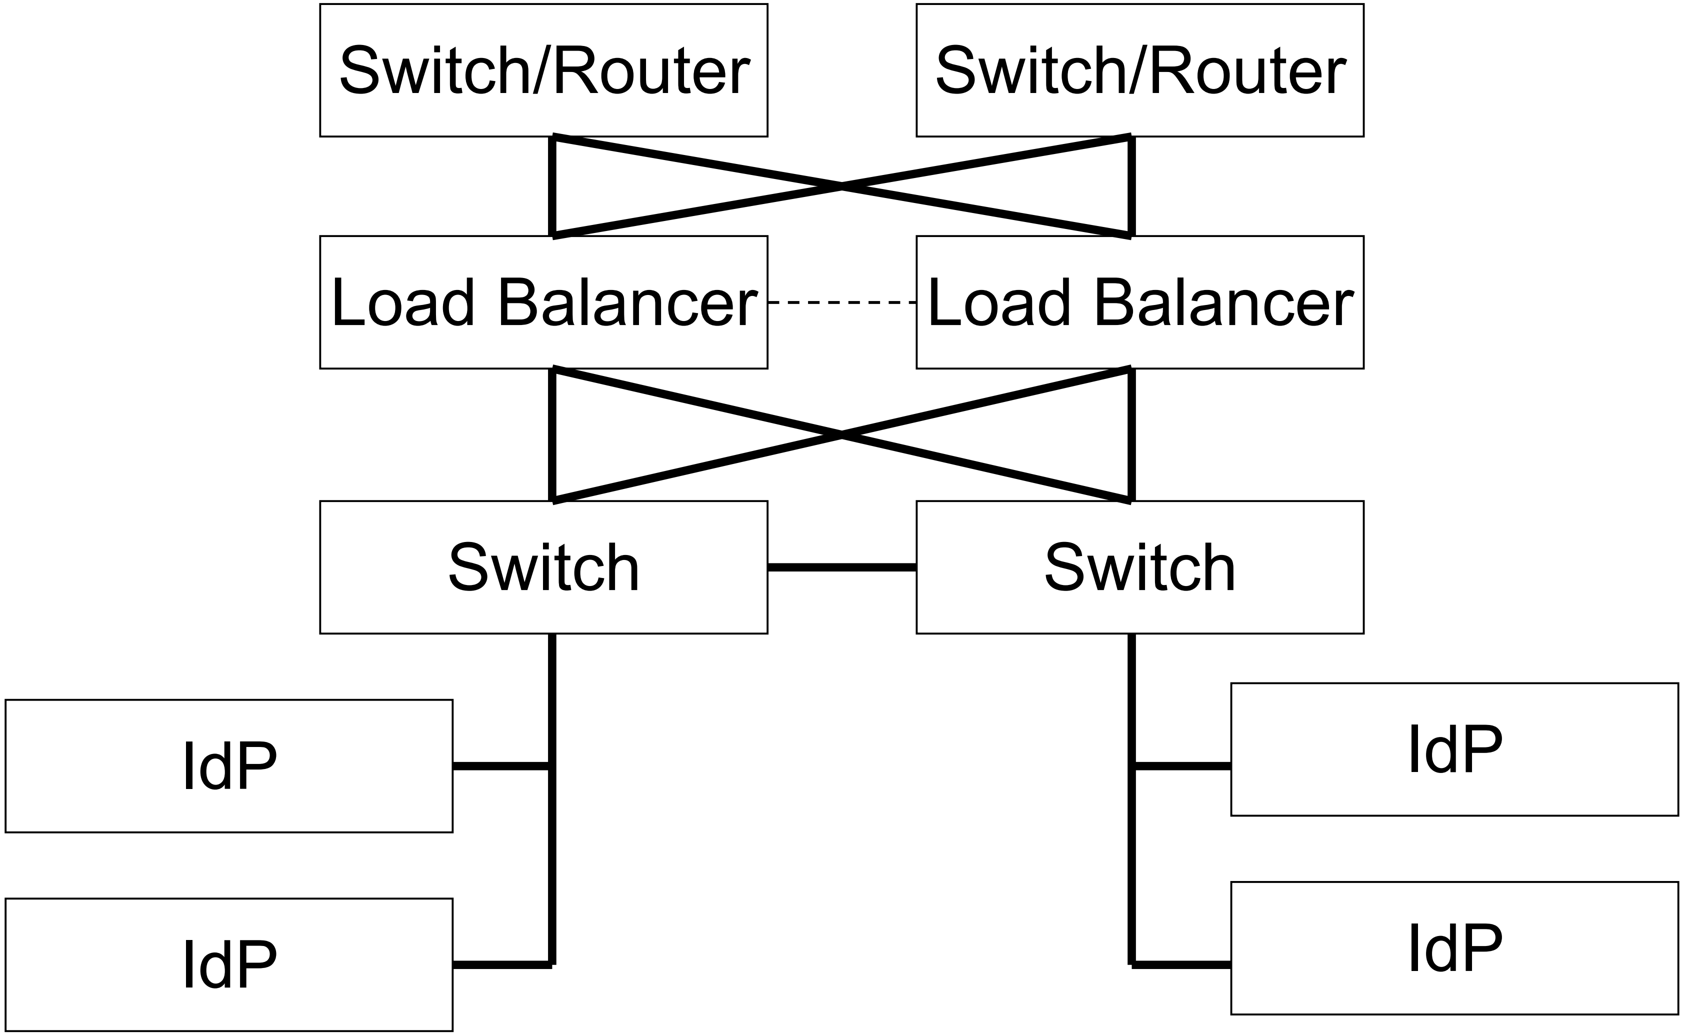

IdP Tuning

- Tune your JVM
- Tune your servlet container
- Within the resolver:
  - Explicitly enumerate the attributes you need from the database and LDAP directory
  - Mark any attribute definition that is never going to be released as dependency only

IdP Configuration

- Do **NOT** enabled configuration reloading, stop and start the IdP and watch the logs for errors.
- Do **NOT** edit configuration files on the production IdP, instead edit and test them in pre-production and then check them into a version control system. Check them out on the production system.

IdP Configuration

- Ensure that your credentials are collected over HTTPS (you may even want to turn off HTTP)
- Ensure you understand what protocols are enabled on the IdP and what their settings are.
- Ensure all your connections to authentication and data services are secure.

IdP Configuration

- Reload your metadata often

IdP Clustering

- IdP already has support for clustering built-in, I just need to document it.
- The IdP will use Terracotta to replicate the JVM heap between cluster elements.
  - This eliminates the need for all replicated data to be serializable.
- Terracotta plugs in to the JVM through the JVMTI.
- Deploy the Terracotta servers on your IdP servers, either in active/active or active/standby mode.

Upcoming Productionalization Related Features

- Two new configuration resources:
  - FileBackedHTTP which can load a configuration file from a central HTTP store
  - SVN which can load a configuration file from a subversion repository
- Configuration resource filters:
  - Allows a resource to be run through one, or more, filters before being loaded
  - Will ship with one implementation that replaces macros with values from a property file